

CLAIMS

What is claimed is:

1. A computer system, comprising:

a pipelined, simultaneous and redundantly threaded ("SRT") processor having at least two threads;

an input/output ("I/O") controller coupled to said processor;

an I/O device coupled to said I/O controller; and

a system memory coupled to said processor;

wherein said SRT processor further comprises:

a load/store execution unit having a store queue, the store queue adapted to store memory requests submitted by the at least two threads, where said memory requests change values in system memory directly or indirectly;

a compare logic coupled to said load/store execution unit;

wherein said compare logic scans the contents of said store queue for corresponding memory requests, and said compare logic verifies that each corresponding memory request matches; and

wherein said compare logic, based on whether the corresponding memory requests match, performs one of 1) allowing the memory request to execute, and 2) initiating fault recovery.

2. The computer system as defined in claim 1 wherein said memory requests that directly or indirectly change data values in system memory further comprise at least committed store requests.

1 8. A method of detecting transient faults in a simultaneously and redundantly threaded
2 microprocessor having at least two threads, the method comprising:
3 executing a program as a first thread;
4 generating a first committed store request from said first thread;
5 storing said first committed store request in a storage queue;
6 executing the program as a second thread;
7 generating a second committed store request from said second thread;
8 storing said second committed store in said storage queue;
9 checking an address and data associated with said first committed store request against an
10 address and data associated with said second committed store request in a compare logic; and
11 allowing one of said first and second committed store requests to execute if the checking
12 step shows those committed store requests are exactly the same.

1 9. The method as defined in claim 8 wherein executing the first and second threads further
2 comprises executing the first thread a plurality of program steps ahead of the second thread.

1 10. The method as defined in claim 9 further comprising:
2 allowing the first and second threads to make speculative branch execution independent of
3 each other.

1 11. The method as defined in claim 9 further comprising:
2 allowing the first thread to execute program steps out of an order of the program;

allowing the second thread to execute program steps out of the order of the program; and

allowing each of the first and second threads to execute the program in a different order from each other.

12. A simultaneous and redundantly threaded microprocessor comprising:

a first pipeline executing a first program thread;

a second pipeline executing a second program thread;

a store queue coupled to each of said first and second pipelines;

a compare circuit coupled to said store queue;

wherein each of said first and second program threads independently generate corresponding committed write requests, and each thread places those committed write requests in the store queue; and

wherein said compare circuit detects transient faults in operation of said first and second pipeline by comparing at least the committed store requests from each thread.

13. A pipelined, simultaneous and redundantly threaded (“SRT”) processor, comprising:

a fetch unit that fetches instructions from a plurality of threads of instructions;

an instruction cache coupled to said fetch unit and storing instructions to be decoded and executed; and

decode logic coupled to said instruction cache to decode the type of instructions stored in
said instruction cache;

wherein said processor processes a set of instructions in a leading thread and also in a trailing thread, and wherein the instructions in the trailing thread are substantially identical to the

1 instructions in the leading thread, the instructions in the trailing thread beginning processing
2 through the processor after the corresponding instructions in the leading thread begin processing
3 through the processor;

4 and wherein said processor detects transient faults by verifying as between the leading and
5 trailing threads only the committed stores and uncached memory read requests.

1 14. A method of detecting transient faults in a simultaneous and redundantly threaded
2 microprocessor having at least two threads, the method comprising:

3 executing a program as a first thread;
4 generating a first committed store request from said first thread;
5 storing said first committed store request in a storage queue;
6 executing the program as a second thread;
7 generating a second committed store request from said second thread;
8 checking an address and data associated with said first committed store request against an
9 address and data associated with said second committed store request; and
10 allowing one of said first and second committed store requests to execute if the checking
11 step shows those committed store requests are exactly the same.

1 15. The method as defined in claim 14 wherein executing the first and second threads further
2 comprises executing the first thread a plurality of program steps ahead of the second thread.

1 16. The method as defined in claim 15 further comprising:

allowing the first and second threads to make speculative branch execution independent of each other.

17. The method as defined in claim 15 further comprising:

allowing the first thread to execute program steps out of an order of the program;
allowing the second thread to execute program steps out of the order of the program; and
allowing each of the first and second threads to execute the program in a different order from each other.

18. A simultaneous and redundantly threaded microprocessor comprising:

a first pipeline executing a first program thread;
a second pipeline executing a second program thread;
a store queue coupled to at least said first pipelines;
wherein each of said first and second program threads independently generate corresponding committed write requests, at least said first thread places the committed write requests in the store queue; and
wherein said second thread detects transient faults in operation of said first and second pipeline by comparing at least the committed store requests from each thread.